

CRUNCHING MORTALITY AND LIFE INSURANCE PORTFOLIOS WITH EXTENDED CREDITRISK⁺

JONAS HIRZ, UWE SCHMOCK, AND PAVEL V. SHEVCHENKO

ABSTRACT. This paper provides a summary of actuarial applications in life insurance of the collective risk model extended CreditRisk⁺ amongst which we find stochastic mortality modelling, joint modelling of underlying death causes as well as profit and loss modelling of life insurance and annuity portfolios as required in (partial) internal models. This approach provides an efficient, numerically stable algorithm for an exact calculation of the one-period profit and loss distribution where various sources of risk can be considered. Model parameters can be estimated by likelihood optimization and Bayesian approaches, using Markov-Chain Monte-Carlo techniques. We provide real world examples using Austrian and Australian death data.

1. INTRODUCTION

Pricing of retirement income products, using internal models and calculating P&L attributions in corresponding portfolios depend crucially on the accuracy of predicted death probabilities. Life insurers and pension funds typically use deterministic generation life tables obtained from crude death rates and then apply some forecasting model or generation life tables. Afterwards, artificial risk margins are often added to account for phenomena associated with longevity, size of the portfolio, selection phenomena, estimation and various other sources. These approaches often lack a stochastic foundation and are certainly not consistently appropriate for all companies due to a possibly twisted mix of these sources of risk. Moreover, we have observed drastic shifts in death rates (yearly deaths divided by population at June 30 in the corresponding year) due to certain underlying death causes over the past decades. An underlying death cause is to be understood as the disease or injury that initiated the train of morbid events leading directly to death. As an illustration of this fact, Figure 1.1 shows death rates based on Australian data for death causes, such as mental and behavioural disorders and circulatory diseases, from 1987 to 2011 for various age categories and both genders. Diseases of the circulatory system, such as ischaemic heart disease, have been clearly reduced throughout the past years while death rates due to mental and behavioural disorders, such as dementia, have doubled for older age groups. This observation nicely illustrates the existence of serial dependence amongst different death causes.

Date: November 28, 2016.

J. Hirz gratefully acknowledges financial support from the Australian Government via the 2014 Endeavour Research Fellowship, as well as from the Oesterreichische Nationalbank (Anniversary Fund, project number: 14977) and Arithmetica. P. V. Shevchenko gratefully acknowledges financial support by the CSIRO-Monash Superannuation Research Cluster, a collaboration among CSIRO, Monash University, Griffith University, the University of Western Australia, the University of Warwick, and stakeholders of the retirement system in the interest of better outcomes for all. Furthermore, we thank Anselm Fleischmann (BELTIOS GmbH, Vienna) for fruitful discussions.

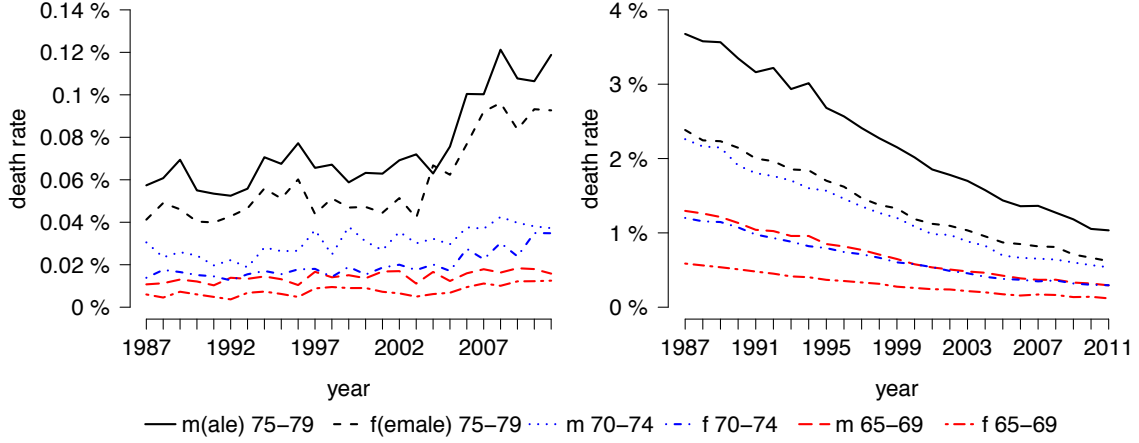


FIGURE 1.1. Australian death rates for mental and behavioural disorders (left) and for circulatory diseases (right).

In our recent work Hirz, Schmock and Shevchenko [2], we are introducing a novel actuarial framework using a collective risk model called extended CreditRisk^+ . In light of the issues described above, it provides a unified and stochastically sound approach to model mortality on the one hand and risk aggregation in annuity as well as life insurance portfolios on the other hand. The general form of this credit risk model was introduced by Schmock [5] and is very different from usual time series approaches such as Lee–Carter or cohort models, see Lee and Carter [4], as well as Cairns et al. [1]. Stochastic modelling of mortality has become increasingly important also because new regulatory requirements such as Solvency II allow internal stochastic models for a more risk-sensitive evaluation of capital requirements and the ability to derive accurate P&L attributions. Our model allows multiple applications, including stochastic modelling of mortality, joint forecasting of death cause intensities, calculation of P&L in life insurance portfolios exactly via Panjer recursion instead of a Monte Carlo method, and partial internal model applications in the (biometric) underwriting risk module. In particular, our model is able to quantify the risk of statistical fluctuations within the next period (experience variance) and parameter uncertainty (change in assumptions).

In this paper we focus on concepts and application aspects of the proposed framework while for details of the proofs and precise mathematical conditions we refer to Hirz, Schmock and Shevchenko [2]. We start with a brief model description in Section 2 and recall different estimation procedures in Section 3. Then, in Section 4, we apply our model and estimation procedures to Australian and Austrian data and give some further applications including mortality forecasts, scenario analysis and a potential internal model application. In Section 5 we give some validation techniques and end with concluding remarks in Section 6.

2. MODEL

Based on the collective risk model *extended CreditRisk⁺*, see Schmock [5], we introduce our model as follows. Let $\{1, \dots, m\}$ denote the set of people (we call them *policyholders* in light of potential applications to annuity portfolios) in the portfolio and let random death indicators N_1, \dots, N_m indicate the *number of deaths* (this may also include or be lapse) of each policyholder in the following period. The event $\{N_i = 0\}$ indicates survival of person i whilst $\{N_i \geq 1\}$ indicates death. Thus, we define death probability $q_i := \mathbb{E}[N_i]$. In reality, death indicators are Bernoulli random variables. Unfortunately in practice, such an approach

is not tractable for calculating loss distributions of large portfolios as execution times explode. Instead, we will assume a mixed Poisson distribution for each policyholder which is very accurate (as can be shown in numerical examples) and which provides an efficient way for calculating loss distributions.

Consider independent *portfolio quantities* Y_1, \dots, Y_m with $d \geq 1$ dimensions which are assumed to be independent of N_1, \dots, N_m . We are then interested in (random) cumulative quantities due to deaths

$$S := \sum_{i=1}^m \sum_{j=1}^{N_i} Y_{i,j},$$

where $(Y_{i,j})_{j \in \mathbb{N}}$ for every i is an i.i.d. sequence of random variables with the same distribution as Y_i . Thus, depending on the specific application of our proposed model, portfolio quantities and the sum S can be interpreted differently:

- (a) For applications in the context of internal models we may set Y_i as the best estimate liability (BEL) of the contract of i such that S gives the sum of BELs of all policyholders who do not survive the following period.
- (b) In the context of portfolio payment analysis we may set Y_i as the payments (such as annuities) to i over the next period such that S gives the sum of payments to all policyholders who do not survive the following period. We may include premiums in a second dimension in order to get joint distributions of premiums and payments.
- (c) For applications in the context of mortality and population estimation we may set $Y_i = 1$ such that S gives the number of deaths in the next period.

In the context of cumulative payments L in an annuity portfolio, for example, it is necessary to consider the sum the payments if everyone survives minus all payments which need not be paid due to deaths, i.e. $L := \sum_{i=1}^m X_i - S$, where X_i may be equal to Y_i or some (random) fraction of sub-annual payments are considered.

To ensure multi-level dependence, we introduce stochastic risk factors $\Lambda_1, \dots, \Lambda_K$ which are independent and gamma distributed with mean one and variance $\sigma_k^2 > 0$. These risk factors can be seen as latent random factors which jointly influence deaths of all policyholders, such as underlying death causes with the stochasticity lying in changes in treatments, better medication, or outbreaks of epidemics, for example. Both the assumption of independence amongst risk factors and the assumption of gamma distributions can be relaxed, see Schmock [5]. Constants $w_{i,0}, \dots, w_{i,K} \in [0, 1]$ for policyholder i represent weights with $w_{i,0} + \dots + w_{i,K} = 1$, indicating the vulnerability of policyholder i to risk factors. For every policyholder i , the total number of deaths N_i is split up according to risk factors as $N_i = N_{i,0} + \dots + N_{i,K}$, where the idiosyncratic $N_{1,0}, \dots, N_{m,0}$ are independent from one another, as well as all other random variables, and Poisson distributed with intensity $q_i w_{i,0}$. Furthermore, conditional on the risk factors, death indicators $\{N_{i,k}\}_{i \in \{1, \dots, m\}, k \in \{1, \dots, K\}}$ are independent and Poisson distributed with random intensity $q_i w_{i,k} \Lambda_k$. Common stochastic risk factors introduce dependence such that $\text{Cov}(N_i, N_j) = q_i q_j \sum_{k=1}^K w_{i,k} w_{j,k} \sigma_k^2$, for all $i \neq j$.

There exists a numerically stable algorithm to derive the distribution of S very efficiently. It can be found in the detailed paper of Hirz, Schmock and Shevchenko [2], or in the lecture notes of Schmock [5, Section 6.7]. Basically, this algorithm uses iterated Panjer recursion to derive $\mathbb{P}(S = n)$ up to every desired $n \in \mathbb{N}$. Approximations arise from mixed Poisson instead of Bernoulli distributions for deaths and—if required—due to stochastic rounding to work with greater loss units. Nevertheless, implementations of this algorithm are significantly

faster than Monte Carlo approximations for comparable error levels. As an illustration we

TABLE 2.1. Quantiles, execution times (speed) and total variation distance (accuracy) of Monte Carlo with 50 000 simulations, as well as our annuity model given a simple portfolio.

	quantiles					speed	accuracy
	1%	10%	50%	90%	99%		
Monte Carlo, $w_{i,0} = 1$	450	472	500	528	552	22.99 sec.	0.0187
annuity model, $w_{i,0} = 1$	449	471	500	529	553	0.01 sec.	0.0125
Monte Carlo, $w_{i,1} = 1$	202	310	483	711	936	23.07 sec.	0.0489
annuity model, $w_{i,1} = 1$	204	309	483	712	944	0.02 sec.	≤ 0.0500

take a portfolio with $m = 10\,000$ policyholders having death probabilities $q := q_i = 0.05$ and payments $Y_i = 1$. We then derive the distribution of S using our annuity model for the case with just idiosyncratic risk, i.e., $w_{i,0} = 1$, and for the case with just one common stochastic risk factor Λ_1 with variance $\sigma_1 = 0.1$ and no idiosyncratic risk, i.e., $w_{i,1} = 1$. Then, using 50 000 simulations of the corresponding model where N_i is Bernoulli distributed or mixed Bernoulli distributed given truncated risk factor $\Lambda_1 | \Lambda_1 \leq \frac{1}{q}$, we compare the results of our annuity model to Monte Carlo, respectively. Truncation of risk factors in the Bernoulli model is necessary as otherwise death probabilities may exceed one. We observe that our annuity model drastically reduces execution times at comparable error levels. Error levels in the purely idiosyncratic case are measured in terms of total variation distance between approximations and the binomial distribution with parameters $(10\,000, 0.05)$ which arises as the independent sum of all Bernoulli random variables. Error levels in the purely non-idiosyncratic case are measured in terms of total variation distance between approximations and the mixed binomial distribution where for our annuity model we use Poisson approximation to get an upper bound. Results are summarised in Table 2.1.

To account for trends in death probabilities and shifts in death causes, we introduce the following parameter families. First, let F^{Lap} denote the cumulative Laplace distribution function with mean zero and variance two, i.e.,

$$F^{\text{Lap}}(x) = \frac{1}{2} + \frac{1}{2} \text{sign}(x) (1 - \exp(-|x|)), \quad x \in \mathbb{R}, \quad (2.1)$$

and trend acceleration and trend reduction with parameters $(\zeta, \eta) \in \mathbb{R} \times (0, \infty)$ is given by $\mathcal{T}_{\zeta, \eta}(t) = (\mathcal{T}_{\zeta, \eta}^*(t) - \mathcal{T}_{\zeta, \eta}^*(t_0)) / (\mathcal{T}_{\zeta, \eta}^*(t_0) - \mathcal{T}_{\zeta, \eta}^*(t_0 - 1))$ with normalisation parameter $t_0 \in \mathbb{R}$ and

$$\mathcal{T}_{\zeta, \eta}(t) = \frac{1}{\eta} \arctan(\eta(t - \zeta)), \quad t \in \mathbb{R}. \quad (2.2)$$

The normalisation above guarantees that parameter β can be compared across different data sets. Given $x < 0$, (2.1) simplifies to an exponential function. Then, death probabilities for all policyholders i , with year of birth $z_i \in \mathbb{R}$, is given by

$$q_i(t) = F^{\text{Lap}}(\alpha_i + \beta_i \mathcal{T}_{\zeta_i, \eta_i}(t) + \kappa_{z_i}), \quad (2.3)$$

where $\alpha_i, \beta_i, \zeta_i, \kappa_{z_i} \in \mathbb{R}$ and $\eta_i \in (0, \infty)$, as well as, for $k \in \{0, \dots, K\}$, weights are given by

$$w_{i,k}(t) = \frac{\exp(u_{i,k} + v_{i,k} \mathcal{T}_{\phi_k, \psi_k}(t))}{\sum_{j=0}^K \exp(u_{i,j} + v_{i,j} \mathcal{T}_{\phi_j, \psi_j}(t))}, \quad (2.4)$$

with $u_{i,0}, v_{i,0}, \phi_0, \dots, u_{i,K}, v_{i,K}, \phi_K \in \mathbb{R}$, as well as $\psi_0, \dots, \psi_K \in (0, \infty)$. Thus, in practical situations, this yields an exponential evolution of death probabilities in (2.3) modulo trend reduction and cohort effects. Expression (2.2) is used for a trend reduction technique which is motivated by Kainhofer, Predota and Schmock [3, Section 4.6.2] and ensures that weights and death probabilities are limited as $t \rightarrow \infty$. Conceptionally, the parameter η gives the speed of trend reduction and the parameter ζ on the other hand gives the shift on the S-shaped arctangent curve, i.e., the location of trend acceleration and trend reduction. Parameter κ models cohort effects for groups with the same or a similar year of birth. This factor can also be understood in a more general context, in the sense that a cohort effect may model categorical variates such as smoker/non-smoker, diabetic/non-diabetic or country of residence. Moreover, cohort effects could be used for modelling weights $w_{i,k}(t)$ but is avoided here as sparse data does not allow proper estimation. Even more restrictive, in applications we often fix ϕ and ψ to reduce dimensionality to suitable levels. Note that in our model every other deterministic trend family can be assumed.

3. ESTIMATION

Considering discrete-time periods $U := \{1, \dots, T\}$ with time index $t \in U$, we assume that age- and calendar year-dependent death probabilities $q_i(t)$ and corresponding weights $w_{i,k}(t)$, see (2.3) and (2.4), are the same for all representative policyholders $i \in \{1, \dots, m\}$ within the same age category $a \in \{1, \dots, A\}$, same gender $g \in \{f, m\}$ and with respect to the same risk factor $\Lambda_k(t)$ with death cause $k \in \{0, \dots, K\}$. For notational purposes we may therefore define $q_{a,g}(t) := q_i(t)$ and $w_{a,g,k}(t) := w_{i,k}(t)$ for a representative policyholder i of age category a and gender g with respect to risk factor $\Lambda_k(t)$, with birth years being denoted by $z_{a,t}$. All random variables at time $t \in \{1, \dots, T\}$ are assumed to be independent of random variables at some different point in time $s \neq t$ with $s \in \{1, \dots, T\}$, as well as risk factors $\Lambda_k(1), \dots, \Lambda_k(T)$ are identically distributed for each k . Given historical population counts $m_{a,g}(t)$ and historical number of deaths $n_{a,g,k}(t)$ due to underlying death causes $k = 0, \dots, K$ we can then derive various estimation procedures if $n_{a,g,k}(t)$ is assumed to be a realisation of the random variable

$$N_{a,g,k}(t) := \sum_{i \in M_{a,g}(t)} N_{i,k}(t),$$

where $M_{a,g}(t) \subset \{1, \dots, m(t)\}$ with $|M_{a,g}(t)| = m_{a,g}(t)$ denotes the set of policyholders of specified age group and gender. Then, the likelihood function $\ell(n|\theta_q, \theta_w, \sigma)$ of parameters $\theta_q := (\alpha, \beta, \zeta, \eta, \kappa)$, as well as $\theta_w := (u, v, \phi, \psi)$ and $\sigma := (\sigma_k)$ given data $n := (n_{a,g,k}(t))$ is given by

$$\begin{aligned} \ell(n|\theta_q, \theta_w, \sigma) &= \prod_{t=1}^T \left(\left(\prod_{a=1}^A \prod_{g \in \{f, m\}} \frac{e^{-\rho_{a,g,0}(t)} \rho_{a,g,0}(t)^{n_{a,g,0}(t)}}}{n_{a,g,0}(t)!} \right) \right. \\ &\quad \times \left. \prod_{k=1}^K \left(\frac{\Gamma(1/\sigma_k^2 + n_k(t))}{\Gamma(1/\sigma_k^2) (\sigma_k^2)^{1/\sigma_k^2} (1/\sigma_k^2 + \rho_k(t))^{1/\sigma_k^2 + n_k(t)}} \prod_{a=1}^A \prod_{g \in \{f, m\}} \frac{\rho_{a,g,k}(t)^{n_{a,g,k}(t)}}{n_{a,g,k}(t)!} \right) \right). \end{aligned} \quad (3.1)$$

where $n_k(t) := \sum_{a=1}^A \sum_{g \in \{f, m\}} n_{a,g,k}(t)$, as well as $\rho_{a,g,k}(t) := m_{a,g}(t) q_{a,g}(t) w_{a,g,k}(t)$ and $\rho_k(t) := \sum_{a=1}^A \sum_{g \in \{f, m\}} \rho_{a,g,k}(t)$.

Since the products in (3.1) can become small, we recommend to use the log-likelihood function instead. Examples suggest that maximum-likelihood estimates are unique. However,

deterministic numerical optimisation routines easily break down due to high dimensionality. Switching to a Bayesian setting, it is straightforward to apply Markov chain Monte Carlo (MCMC) methods to get many samples $\theta = (\theta_q^i, \theta_w^i, \sigma^i)$ from the posterior distribution $\pi(\theta|n) \propto \ell(n|\theta)\pi(\theta)$ where $\pi(\theta)$ denotes the prior distribution of parameters. The mean over these samples provides a good point estimate for model parameters and sampled posterior distribution can be used to estimate parameter uncertainty. Various MCMC algorithms are available. Hirz, Schmock and Shevchenko [2] implement a well-known random walk Metropolis–Hastings within Gibbs algorithm, in which case the mode of the posterior samples corresponds to a maximum-likelihood estimate. The method requires a certain burn-in period until the generated chain becomes stationary. To reduce long computational times, one can run several independent MCMC chains with different starting points on different CPUs in a parallel way. To prevent overfitting, it is possible to regularise, i.e., smooth, maximum a posteriori estimates via adjusting the prior distribution $\pi(\theta)$. This technique is particularly used in regression, as well as in many applications, such as signal processing. When forecasting death probabilities in Section 4.4, we use a Gaussian prior distribution with a certain correlation structure.

If risk factors are not integrated out, under a Bayesian setting, we may also derive the posterior distribution of the risk factors. Necessarily in that case, estimates $\hat{\lambda}$ for risk factors and $\hat{\sigma}$ for risk factor variances satisfy

$$\hat{\lambda}_k^{\text{MAP}}(t) = \frac{1/(\hat{\sigma}_k^{\text{MAP}})^2 - 1 + \sum_{a=1}^A \sum_{g \in \{f,m\}} n_{a,g,k}(t)}{1/(\hat{\sigma}_k^{\text{MAP}})^2 + \sum_{a=1}^A \sum_{g \in \{f,m\}} \rho_{a,g,k}(t)} \quad (3.2)$$

if $1/(\hat{\sigma}_k^{\text{MAP}})^2 - 1 + \sum_{a=1}^A \sum_{g \in \{f,m\}} n_{a,g,k}(t) > 0$, as well as

$$2 \log \hat{\sigma}_k^{\text{MAP}} + \frac{\Gamma'(1/(\hat{\sigma}_k^{\text{MAP}})^2)}{\Gamma(1/(\hat{\sigma}_k^{\text{MAP}})^2)} = \frac{1}{T} \sum_{t=1}^T (1 + \log \hat{\lambda}_k^{\text{MAP}}(t) - \hat{\lambda}_k^{\text{MAP}}(t)), \quad (3.3)$$

where, for given $\hat{\lambda}_k^{\text{MAP}}(1), \dots, \hat{\lambda}_k^{\text{MAP}}(T) > 0$, (3.3) has a unique solution which is strictly positive. Thus, rougher—but still very accurate—estimates are given by

$$\hat{\lambda}_k^{\text{MAPappr}}(t) := \frac{-1 + \sum_{a=1}^A \sum_{g \in \{f,m\}} n_{a,g,k}(t)}{\sum_{a=1}^A \sum_{g \in \{f,m\}} \rho_{a,g,k}(t)} \quad (3.4)$$

as well as

$$\hat{\sigma}_k^{\text{MAPappr}} := \sqrt{\frac{1}{T} \sum_{t=1}^T (\hat{\lambda}_k^{\text{MAPappr}}(t) - 1)^2}, \quad k \in \{1, \dots, K\},$$

which is simply the sample variance of $\hat{\lambda}^{\text{MAP}}$.

Much faster to derive but also less accurate, we can use a matching of moments approach. Therefore, we have to make the simplifying assumption that deaths $(N_{a,g,k}(t))_{t \in U}$ are i.i.d. which can approximatively be obtained by modifying the deaths $n_{a,g,k}(t)$ via

$$n'_{a,g,k}(t) := \left\lfloor \frac{m_{a,g}(T) q_{a,g}(T) w_{a,g,k}(T)}{m_{a,g}(t) q_{a,g}(t) w_{a,g,k}(t)} n_{a,g,k}(t) \right\rfloor, \quad t \in U, \quad (3.5)$$

and, correspondingly, $m_{a,g} := m_{a,g}(T)$, as well as $q_{a,g} := q_{a,g}(T)$ and $w_{a,g,k} := w_{a,g,k}(T)$.

Estimates $\hat{q}_{a,g}^{\text{MM}}(t)$ for death probabilities $q_{a,g}(t)$ can be obtained via minimising mean squared error to death rates which, if parameters ζ , η and κ are previously fixed, can be

obtained by regressing

$$(F^{\text{Lap}})^{-1} \left(\frac{\sum_{k=0}^K n'_{a,g,k}(t)}{m_{a,g}(t)} \right) - \kappa_{z_{a,g}}$$

on $\mathcal{T}_{\zeta_{a,g}, \eta_{a,g}}(t)$. Estimates $\hat{u}_{a,g,k}^{\text{MM}}, \hat{v}_{a,g,k}^{\text{MM}}, \hat{\phi}_k^{\text{MM}}, \hat{\psi}_k^{\text{MM}}$ for parameters $u_{a,g,k}, v_{a,g,k}, \phi_k, \psi_k$ via minimising the mean squared error to death rates which again, if parameters ϕ and ψ are previously fixed, can be obtained by regressing

$$\log \frac{n'_{a,g,k}(t)}{m_{a,g}(t) \hat{q}_{a,g}^{\text{MM}}(t)}$$

on $\mathcal{T}_{\phi_k, \psi_k}(t)$. Estimates $\hat{w}_{a,g,k}^{\text{MM}}(t)$ are then given by (2.4).

Defining unbiased estimators for weights $W_{a,g,k}^*(t) := N_{a,g,k}(t)/(m_{a,g} q_{a,g})$, as well as $\bar{W}_{a,g,k}^* := \frac{1}{T} \sum_{t=1}^T W_{a,g,k}^*(t)$ gives estimator

$$\hat{\Sigma}_{a,g,k}^2 = \frac{1}{T-1} \sum_{t=1}^T (W_{a,g,k}^*(t) - \bar{W}_{a,g,k}^*)^2,$$

we have

$$\mathbb{E}[\hat{\Sigma}_{a,g,k}^2] = \text{Var}(W_{a,g,k}^*(t)) = \frac{w_{a,g,k}}{m_{a,g} q_{a,g}} + \sigma_k^2 w_{a,g,k}^2.$$

Thus, *matching of moments estimate* for σ_k can be defined as

$$\hat{\sigma}_k^{\text{MM}} := \sqrt{\max \left\{ 0, \frac{\sum_{a=1}^A \sum_{g \in \{f,m\}} \left(\hat{\sigma}_{a,g,k}^2 - \frac{w_{a,g,k}^{\text{MM}}(T)}{m_{a,g} q_{a,g}^{\text{MM}}(T)} \right)}{\sum_{a=1}^A \sum_{g \in \{f,m\}} (w_{a,g,k}^{\text{MM}}(T))^2} \right\}},$$

where $\hat{\sigma}_{a,g,k}^2$ is the estimate corresponding to estimator $\hat{\Sigma}_{a,g,k}^2$.

4. REAL WORLD EXAMPLE

4.1. Prediction of death cause intensities. As an applied example for estimation in our model, as well as for some further applications, we take annual death data from Australia for the period 1987 to 2011. We fit our annuity model using the matching of moments approach, as well as the maximum-likelihood approach with Markov chain Monte Carlo (MCMC). Data source for historical Australian population, categorised by age and gender, is taken from the Australian Bureau of Statistics and data for the number of deaths categorised by death cause and divided into eight age categories, i.e., 50–54 years, 55–59 years, 60–64 years, 65–69 years, 70–74 years, 75–79 years, 80–84 years and 85+ years, denoted by a_1, \dots, a_8 , respectively, for each gender is taken from the AIHW. The provided death data is divided into 19 different death causes—based on the ICD-9 or ICD-10 classification—where we identify the following ten of them with common non-idiosyncratic risk factors: ‘*certain infectious and parasitic diseases*’, ‘*neoplasms*’, ‘*endocrine, nutritional and metabolic diseases*’, ‘*mental and behavioural disorders*’, ‘*diseases of the nervous system*’, ‘*circulatory diseases*’, ‘*diseases of the respiratory system*’, ‘*diseases of the digestive system*’, ‘*external causes of injury and poisoning*’, ‘*diseases of the genitourinary system*’. We merge the remaining eight death causes to idiosyncratic risk as their individual contributions to overall death counts are small for all categories. Data handling needs some care as there was a change in classification of death data in 1997 as explained at the website of the Australian Bureau of Statistics. Australia introduced the tenth revision of the International Classification of Diseases (ICD-10, following ICD-9) in 1997, with a transition period from 1997 to 1998. Within this period,

comparability factors were produced. Thus, for the period 1987 to 1996, death counts have to be multiplied by corresponding comparability factors.

Trends are considered as described above where cohort and trend reduction parameters are fixed a priori with values $t_0 = 1987$, $\kappa_{z_i} = 0$, $\zeta_{a_i,g} = \phi_k = 0$ and $\eta_{a_i,g} = \psi_k = \frac{1}{150}$ where the last value is based on mean trend reduction in death probabilities with the stated parameter combination. For a more advanced modelling of trend reduction see Figure 4.5. Thus, within the maximum-likelihood framework, we end up with 394 parameters, with 362 to be optimised, in our annuity model. As deterministic numerical optimisation of the likelihood function breaks down due to high dimensionality, we use MCMC in this maximum-likelihood setting instead. Based on 40 000 MCMC steps with burn-in period of 10 000 we are able to derive estimates of all parameters. The execution time of our algorithm is roughly seven hours on a standard computer in ‘R’, several parallel MCMC chains can be run, each with different starting values. We can use Equation (3.2) to derive approximations for risk factor

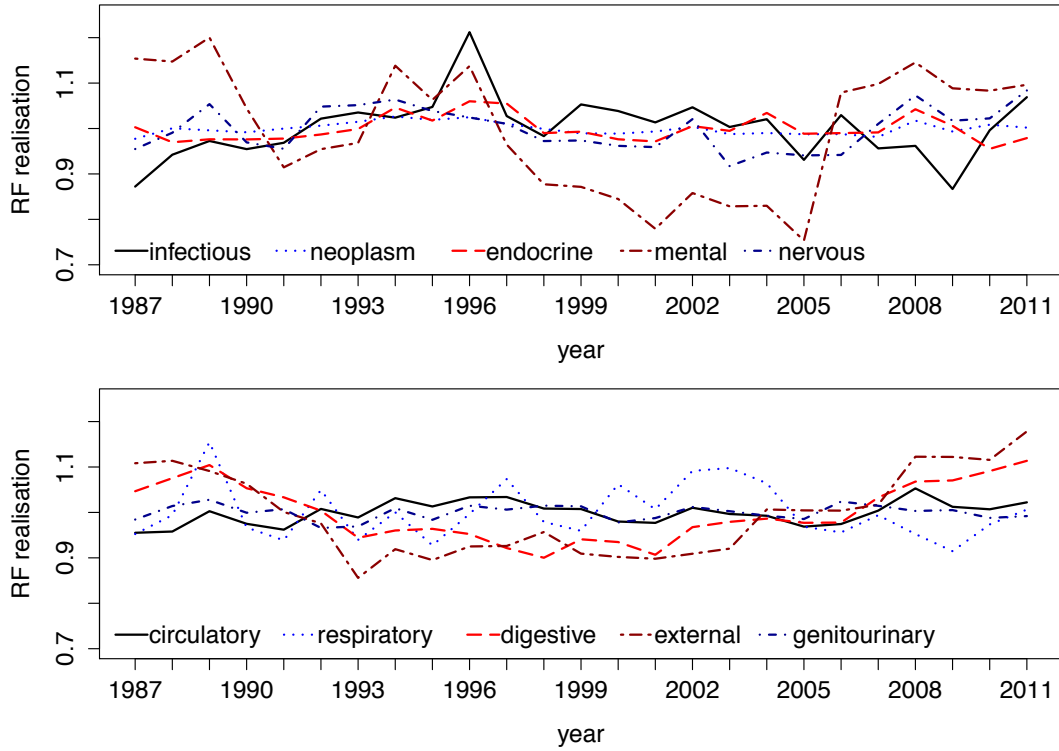


FIGURE 4.1. Estimated risk factor realisations.

realisation estimates where all required parameters are taken from the MCMC estimation. Results are shown in Figure 4.1. In the top figure we observe a massive jump in the risk factor for mental and behavioural disorders between 2005 to 2006 which is mainly driven by an unexpectedly high increase in deaths due to dementia which is explained by the ABS as being due to new coding instructions and changes to various Acts, see Chapter 2.4 of the report on Dementia in Australia of the AIHW (2012).

Assumption (2.4) provides a joint forecast of all death cause intensities, i.e. weights, simultaneously—in contrast to standard procedures where projections are made for each death cause separately. As already presumed in Figure 1.1 in the introduction, our model observes major shifts in weights of certain death causes over previous years as shown in

TABLE 4.1. Estimated weights for all death causes with five and 95 percent quantiles in brackets.

	male				female			
	60 to 64 years		80 to 84 years		60 to 64 years		80 to 84 years	
	2011	2031 (quant.)	2011	2031 (quant.)	2011	2031 (quant.)	2011	2031 (quant.)
neop.	0.499	0.547 (0.561 0.531)	0.324	0.378 (0.392 0.364)	0.592	0.648 (0.662 0.629)	0.263	0.303 (0.319 0.288)
circ.	0.228	0.116 (0.123 0.109)	0.325	0.173 (0.181 0.164)	0.140	0.060 (0.065 0.055)	0.342	0.149 (0.158 0.140)
ext.	0.056	0.062 (0.073 0.053)	0.026	0.028 (0.033 0.024)	0.072	0.069 (0.078 0.060)	0.100	0.126 (0.139 0.113)
resp.	0.051	0.036 (0.040 0.032)	0.106	0.092 (0.101 0.083)	0.038	0.037 (0.043 0.032)	0.051	0.068 (0.074 0.061)
endo.	0.044	0.062 (0.070 0.055)	0.047	0.077 (0.084 0.070)	0.036	0.051 (0.060 0.043)	0.054	0.080 (0.089 0.071)
dig.	0.041	0.036 (0.040 0.031)	0.027	0.020 (0.023 0.018)	0.035	0.032 (0.038 0.026)	0.024	0.023 (0.027 0.020)
nerv.	0.029	0.052 (0.061 0.045)	0.045	0.061 (0.068 0.055)	0.031	0.024 (0.029 0.020)	0.034	0.023 (0.027 0.020)
idio.	0.018	0.028 (0.034 0.023)	0.015	0.018 (0.020 0.016)	0.022	0.023 (0.028 0.019)	0.023	0.024 (0.027 0.022)
inf.	0.014	0.025 (0.033 0.020)	0.015	0.022 (0.027 0.019)	0.014	0.020 (0.027 0.015)	0.017	0.024 (0.028 0.020)
ment.	0.013	0.027 (0.036 0.019)	0.041	0.105 (0.130 0.078)	0.012	0.032 (0.046 0.021)	0.062	0.155 (0.188 0.118)
geni.	0.008	0.008 (0.010 0.006)	0.028	0.025 (0.028 0.023)	0.009	0.005 (0.006 0.004)	0.029	0.026 (0.028 0.023)


Tables 4.1 and 4.2. This table lists weights $w_{a,g,k}(t)$ for all death causes estimated for year 2011, as well as forecasted for 2031 using (2.4) with MCMC mean estimates for ages 60 to 64 years (left) and 80 to 84 years (right). Our model forecasts suggest that if these trends in weight changes persist, then the future gives a whole new picture of mortality. First, deaths due to circulatory diseases are expected to decrease whilst neoplasms will become the leading death cause over most age categories. Moreover, deaths due to mental and behavioural disorders are expected to rise massively for older ages. This observation nicely illustrates the serial dependence, amongst different death causes captured by our model. This potential increase in deaths due to mental and behavioural disorders for older ages will have a massive impact on social systems as, typically, such patients need long-term geriatric care. High uncertainty in forecasted weights is reflected by wide confidence intervals (values in brackets) for the risk factor of mental and behavioural disorders. These confidence intervals are derived from corresponding MCMC chains and, therefore, solely reflect uncertainty associated with parameter estimation. Note that results for estimated trends depend on the length of the data period as short-term trends might not coincide with mid- to long-term trends.

4.2. Scenario analysis. As an explanatory example of our annuity model, assume $m = 1\,600$ policyholders which distribute uniformly over all age categories and genders, i.e., each age category contains 100 policyholders with corresponding death probabilities, as well as weights as previously estimated and forecasted for 2012. Annuities $X_i = Y_i$ for all $i \in \{1, \dots, m\}$ are paid annually and take deterministic values in $\{11, \dots, 20\}$ such that ten policyholders in each age and gender category share equally high payments. We now want to analyse the scenario, indexed by ‘scen’, that deaths due to neoplasms are reduced by 25 percent in 2012 over all age categories. In that case, we can estimate the realisation of risk factor for neoplasms, see (3.2), which takes an estimated value of 0.7991. Running our annuity model with this risk factor realisation being fixed, we end up with a loss distribution L^{scen} where

TABLE 4.2. Leading death causes with weights in brackets.

		male		female	
		2011	2051	2011	2051
55–59 years	1.	neoplasms (0.469)	neoplasms (0.474)	neoplasms (0.603)	neoplasms (0.581)
	2.	circulatory (0.222)	infectious (0.092)	circulatory (0.112)	nervous (0.077)
	3.	external (0.085)	external (0.083)	respiratory (0.058)	not elsewhere (0.068)
65–69 years	1.	neoplasms (0.505)	neoplasms (0.575)	neoplasms (0.551)	neoplasms (0.609)
	2.	circulatory (0.226)	endocrine (0.082)	circulatory (0.162)	mental (0.112)
	3.	respiratory (0.072)	mental (0.075)	respiratory (0.083)	nervous (0.065)
75–79 years	1.	neoplasms (0.405)	neoplasms (0.466)	neoplasms (0.365)	neoplasms (0.378)
	2.	circulatory (0.277)	mental (0.185)	circulatory (0.271)	mental (0.245)
	3.	respiratory (0.100)	endocrine (0.098)	respiratory (0.103)	respiratory (0.108)
85+ years	1.	circulatory (0.395)	mental (0.329)	circulatory (0.441)	mental (0.503)
	2.	neoplasms (0.217)	neoplasms (0.216)	neoplasms (0.131)	circulatory (0.092)
	3.	respiratory (0.115)	circulatory (0.133)	mental (0.101)	neoplasms (0.090)

deaths due to neoplasms have decreased. Figure 4.2 then shows probability distributions of traditional loss L without scenario, as well as of scenario loss L^{neo} with corresponding 95 percent and 99 percent quantiles. We observe that a reduction of 25 percent in cancer death rates leads to a remarkable shift in quantiles of the loss distribution as fewer people die and, thus, more annuity payments have to be made.

FIGURE 4.2. Loss distributions of L and L^{scen} with 95 and 99 percent quantiles.

4.3. Forecasting death rates and comparison with the Lee–Carter model. We can compare out-of-sample forecasts of death rates from our model to forecasts obtained by other mortality models. In this paper, we choose the traditional Lee–Carter model as a proxy, see, for example, Kainhofer, Predota and Schmock [3, Section 4.5.1], as our model is conceptionally based on a similar idea and as it is widely used in practice. Henceforth, we set $Y_j(t) := 1$ for all people $j \in M_{a,g}(t)$ with $M_{a,g}(t) = \{1, \dots, m_{a,g}(T)\}$ and $t \in \{T+1, \dots, S\}$. The simplifying

assumption of a constant population is conservative as population tends to increase and, therefore, statistical fluctuations are over-estimated. To show a further application of our model we compare out-of-sample forecasts from our model to forecasts obtained by the traditional Lee–Carter model. Given the number of living people $m_{a,g}(t)$, as well as annual deaths $n_{a,g}(t) := \sum_{k=0}^K n_{a,g,k}(t)$, the Lee–Carter approach models logarithmic death rates $\log r_{a,g}(t) := \log n_{a,g}(t) - \log m_{a,g}(t)$ in the form $\log r_{a,g}(t) = \mu_{a,g} + \tau_t \nu_{a,g} + \varepsilon_{a,g,t}$ with independent normal error terms $\varepsilon_{a,g,t}$ with mean zero and common time-specific components τ_t . Using suitable normalisations, estimates for these components can be derived via method of moments and singular value decompositions, see, for example, Kainhofer, Predota and Schmock [3, Section 4.5.1]. Forecasts may then be obtained by using auto-regressive models for τ . Conversely, using our model it is straight-forward to *forecast death rates* and to give corresponding confidence intervals via setting $Y_j(t) := 1$ for all people $j \in M_{a,g}(t)$ with $M_{a,g}(t) = \{1, \dots, m_{a,g}(T)\}$ and $t \in \{T+1, \dots, S\}$. The assumption of constant population for forecasts in Australia is conservative as population tends to increase and, therefore, statistical fluctuations are over-estimated. As an alternative, more sophisticated models for population forecasts can be used.

Then, for an estimate $\hat{\theta}$ of parameter vector θ run our annuity model with parameters forecasted, see (2.3) and (2.4). We then obtain the distribution of the total number of deaths $S_{a,g}(t)$ given $\hat{\theta}$ and, thus, forecasted death rate $\hat{r}_{a,g}(t)$ is given by

$$\mathbb{P}\left(\hat{r}_{a,g}(t) = \frac{n}{m_{a,g}(T)}\right) = \mathbb{P}(S_{a,g}(t) = n), \quad n \in \mathbb{N}_0.$$

Uncertainty in the form of confidence intervals represent statistical fluctuations, as well as random changes in risk factors. Additionally, using results obtained by Markov chain Monte Carlo (MCMC) it is even possible to incorporate parameter uncertainty into predictions. To account for an increase in uncertainty for forecasts we suggest to assume increasing risk factor variances for forecasts, e.g., $\tilde{\sigma}_k^2(t) = \sigma_k^2(1 + d(t - T))^2$ with $d \geq 0$. A motivation for this approach with $k = 1$ is the following: A major source of uncertainty for forecasts lies in an unexpected deviation from the estimated trend for death probabilities. We may therefore assume that rather than being deterministic, forecasted values $q_{a,g}(t)$ are beta distributed (now denoted by $Q_{a,g}(t)$) with $\mathbb{E}[Q_{a,g}(t)] = q_{a,g}(t)$ and variance $\sigma_{Q,a,g}^2(t)$ which is increasing in time. Then, given independence amongst risk factor Λ_1 and $Q_{a,g}(t)$, we may assume that there exists a future point t_0 such that

$$\sigma_{Q,a,g}^2(t_0) = \frac{q_{a,g}(t_0)(1 - q_{a,g}(t_0))}{1/\sigma_1^2 + 1}.$$

In that case, $Q_{a,g}(t)\Lambda_1$ is again gamma distributed with mean one and increased variance $q_{a,g}(t_0)\sigma_1^2$ (instead of $q_{a,g}^2(t_0)\sigma_1^2$ for the deterministic case). Henceforth, it seems reasonable to stay within the family of gamma distributions for forecasts and just adapt variances over time. Of course, families for these variances for gamma distributions can be changed arbitrarily and may be selected via classical information criteria.

Using in-sample data, d can be estimated using (3.1) with all other parameters being fixed. Using Australian death and population data for the years 1963 to 1997 we estimate model parameters via MCMC in our annuity model with one common stochastic risk factor having constant weight one. In average, i.e., for various forecasting periods and starting at different dates, parameter d takes the value 0.22 in our example. Using fixed trend parameters as above, and using the mean of 30 000 MCMC samples, we forecast death rates and corresponding confidence intervals out of sample for the period 1998 to 2013. We can

then compare these results to realised death rates within the stated period and to forecasts obtained by the Lee–Carter model which is shown in Figure 4.3 for females aged 50 to 54 years. We observe that true death rates mostly fall in the 90 percent confidence band for both procedures. Moreover, Lee–Carter forecasts lead to wider spreads of quantiles in the future whilst our model suggests a more moderate increase in uncertainty. Taking various parameter samples from the MCMC chain and deriving quantiles for death rates, we can extract contributions of parameter uncertainty in our model coming from posterior distributions of parameters.




FIGURE 4.3. Forecasted death rates using our annuity model (AM) and the Lee–Carter model (LC).

4.4. Forecasting death probabilities. Forecasting death probabilities within our annuity model is straight forward using (2.3). In the special case with just idiosyncratic risk, i.e., $K = 0$, death indicators can be assumed to be Bernoulli distributed instead of being Poisson distributed in which case we may write the likelihood function in the form

$$\ell^B(n|\alpha, \beta, \zeta, \eta, \kappa) = \prod_{t=1}^T \prod_{a=1}^A \prod_{g \in \{f, m\}} \binom{m_{a,g}(t)}{n_{a,g,0}(t)} q_{a,g}(t)^{n_{a,g,0}(t)} (1 - q_{a,g}(t))^{m_{a,g}(t) - n_{a,g,0}(t)},$$

with $0 \leq n_{a,g,0}(t) \leq m_{a,g}(t)$. Due to possible overfitting, derived estimates may not be sufficiently smooth across age categories $a \in \{1, \dots, A\}$. Therefore, if we switch to a Bayesian setting, we may use regularisation via prior distributions. To guarantee smooth results and a sufficient stochastic foundation, we suggest the usage of Gaussian priors with mean zero and a specific correlation structure, i.e., $\pi(\alpha, \beta, \zeta, \eta, \kappa) = \pi(\alpha)\pi(\beta)\pi(\zeta)\pi(\eta)\pi(\kappa)$ with

$$\log \pi(\alpha) := -c_\alpha \sum_{g \in \{f, m\}} \left(\sum_{a=1}^{A-1} (\alpha_{a,g} - \alpha_{a+1,g})^2 + \varepsilon_\alpha \sum_{a=1}^A \alpha_{a,g}^2 \right) + \log(d_\alpha), \quad c_\alpha, d_\alpha, \varepsilon_\alpha > 0, \quad (4.1)$$

and correspondingly for β , ζ , η and κ . Parameters c_α (correspondingly for β , ζ , η and κ) is a scaling parameters and directly associated with the variance of Gaussian priors while

normalisation-parameter d_α guarantees that $\pi(\alpha)$ is a proper Gaussian density. Penalty-parameter ε_α scales the correlation amongst neighbour parameters in the sense that the lower it gets, the higher the correlation. The more we increase c_α the stronger the influence of, or the believe in the prior distribution. This particular prior distribution penalises deviations from the ordinate which is a mild conceptual shortcoming as this does not accurately reflect our prior believes. Nevertheless, it yields good results and allows a direct analysis of prior variances and covariances of parameters. Setting $\varepsilon_\alpha = 0$ gives an improper prior with uniformly distributed (on \mathbb{R}) marginals such that we gain that there is no prior believe in expectations of parameters but, simultaneously, lose the presence of variance-covariance-matrices and asymptotically get perfect positive correlation across parameters of different ages. However, setting $\varepsilon_\alpha = 0$ very often gives better fits of estimates. An optimal choice of regularisation parameters $c_\alpha, c_\beta, c_\zeta, c_\eta$ and c_κ can be obtained by cross-validation. Furthermore, we set $\varepsilon_\alpha = \varepsilon_\beta = 10^{-2}$ and $\varepsilon_\zeta = \varepsilon_\eta = \varepsilon_\kappa = 10^{-4}$ as this yields a suitable prior correlation structure which decreases with higher age differences and which is always positive, see the left plot in Figure 4.4.




FIGURE 4.4. Correlation structure of Gaussian priors with penalisation for deviation from ordinate (left), straight line (middle) and parabola (right).

There exist many other reasonable choices for Gaussian prior distributions. For example, replacing graduation terms $(\alpha_{a,g} - \alpha_{a+1,g})^2$ in (4.1) by higher order differences of the form $(\sum_{\nu=0}^k (-1)^\nu \binom{k}{\nu} \alpha_{a,g+\nu})^2$ yields a penalisation for deviations from a straight line with $k = 2$, see middle plot in Figure 4.4, or from a parabola with $k = 3$, see right plot in Figure 4.4. The usage of higher order differences for graduation of statistical estimates goes back to the Whittaker–Henderson method. Taking $k = 2, 3$ unfortunately yields negative prior correlations amongst certain parameters which is why we do not recommend their use. Of course, there exist many further possible choices for prior distributions.

Results for Australian data from 1971 to 2013 with $t_0 = 2013$ are given in Figure 4.5. Using MCMC we derive estimates for logarithmic death probabilities $\log q_{a,g}(t)$ with corresponding forecasts, mortality trends $\beta_{a,g}$, as well as trend reduction parameters $\zeta_{a,g}, \eta_{a,g}$ and cohort effects $\kappa_{z_{a,g}}$. Recall that $\zeta_{a,g}$ yields the point in time when trend acceleration transits into trend reduction. Values close zero indicate that trend reduction is about to start whilst negative values indicate that trend reduction is already in place. We observe negligible parameter uncertainty due to a long period of data. Further, regularisation parameters obtained by cross-validation are given by $c_\alpha = 1\,300$, $c_\beta = c_\eta = 30\,000c_\alpha$, $c_\zeta = c_\alpha/20$ and $c_\kappa = 1\,000c_\alpha$.

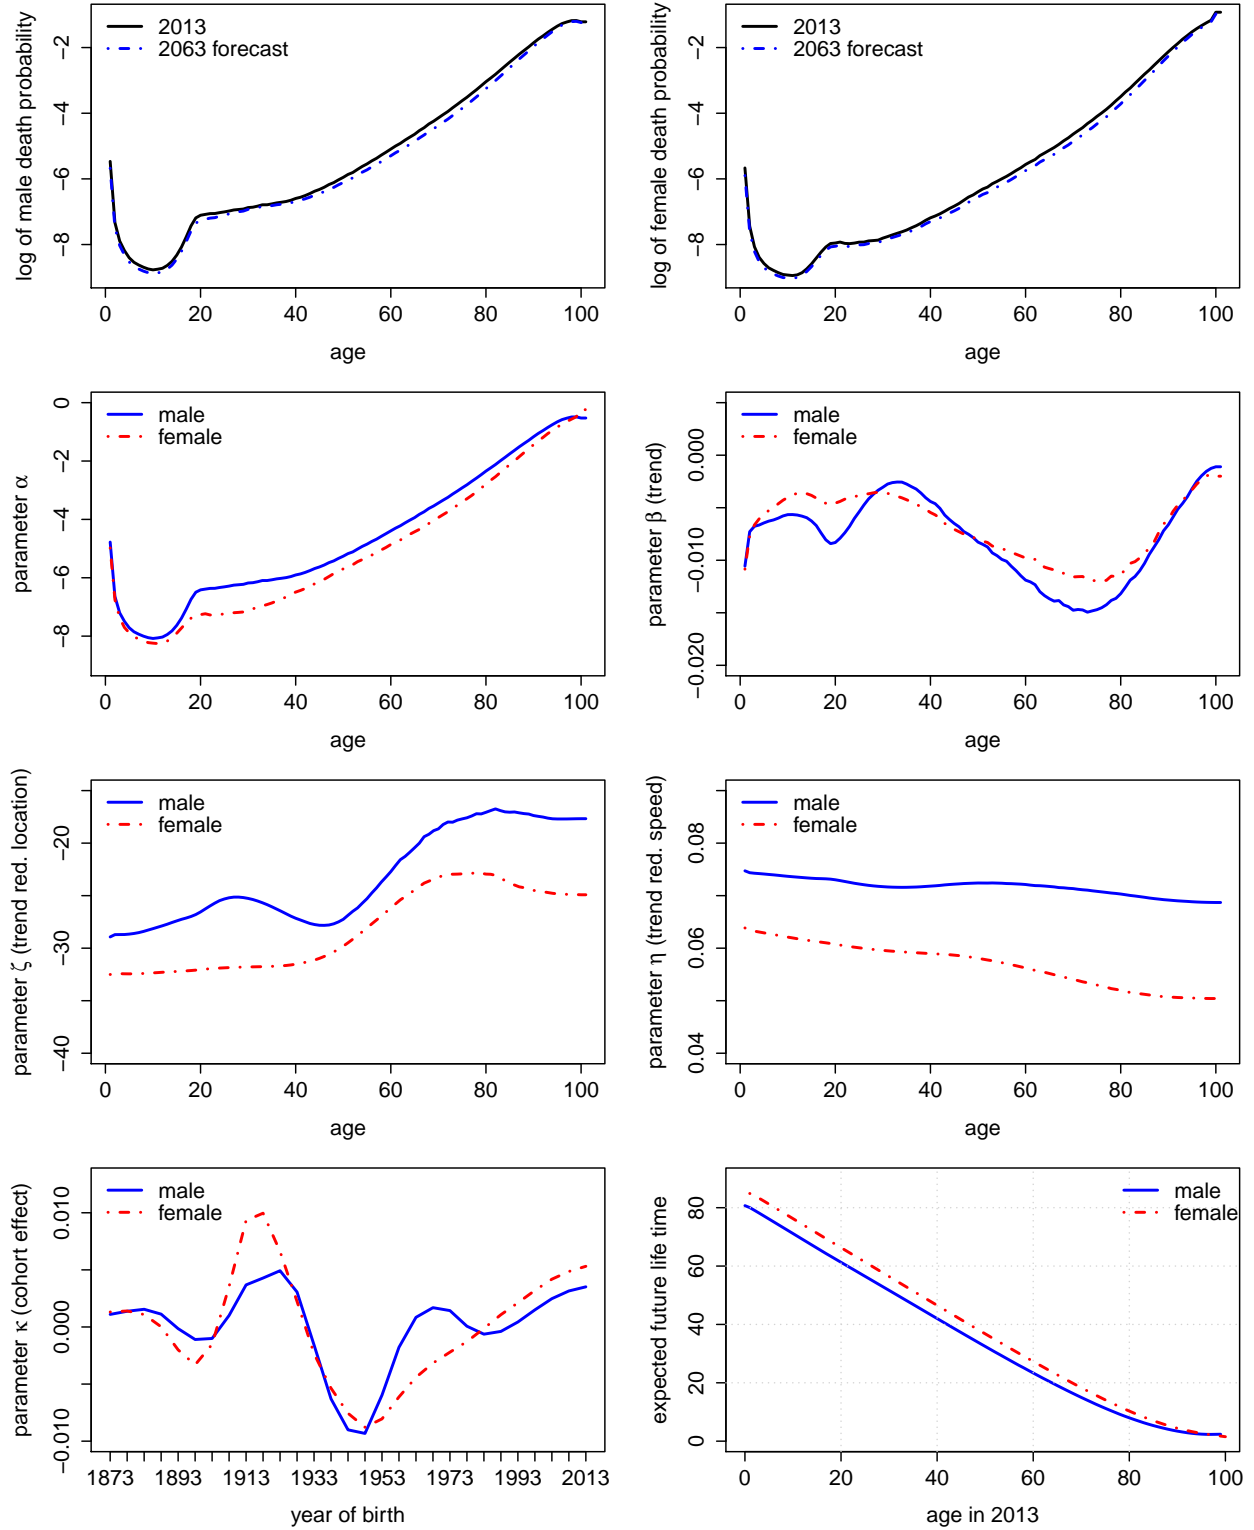


FIGURE 4.5. Logarithm of death probabilities (top) for 2013 in Australia and forecasts for 2063 based on data from 1971 to 2013, as well as corresponding parameters $\alpha, \beta, \zeta, \eta$ and κ , as well as expected future lifetime in 2013.

We can draw some immediate conclusions. Firstly, we see an overall improvement in mortality over all ages where the trend is particularly strong for young ages and ages between

50 and 80 whereas the trend vanishes towards the age of 100, maybe implying a natural barrier for life expectancy. Due to sparse data the latter conclusion should be treated with the utmost caution. Furthermore, we see the classical hump of increased mortality driven by accidents around the age of 20 which is more developed for males.

Secondly, estimates for $\zeta_{a,g}$ suggest that trend acceleration switched to trend reduction throughout the past 20 to 30 years, with a similar shape for males and females. Estimates for $\eta_{a,g}$ show that the speed of trend reduction is unexpectedly high, even stronger for males. It should be noted that estimates for $\zeta_{a,g}$ and $\eta_{a,g}$ are sensitive to penalty-parameters $\varepsilon_\zeta, \varepsilon_\eta$. Estimates for $\kappa_{z_{a,g}}$ show that the cohort effect is particularly strong (in the sense of increased mortality) for the generation born around 1915—probably associated with World War II—and particularly weak for the generation born around 1945.

Henceforth, based on forecasts for death probabilities, expected future life time can be estimated. To be consistent concerning longevity risk, mortality trends have to be included as a 60-year-old today will probably not have as good medication as a 60-year-old in several decades. However, it seems that this is not the standard approach in the literature. Based on the definitions above, expected (curtate) future life time of a person at date T is given by

$$e_{a,g}(T) = \mathbb{E}[K_{a,g}(T)] = \sum_{k=1}^{\infty} {}_k p_{a,g}(T) \quad (4.2)$$

where survival probabilities over $k \in \mathbb{N}$ years are given by ${}_k p_{a,g}(T) := \prod_{j=0}^{k-1} (1 - q_{a+j,g}(T+j))$ and where $K_{a,g}(T)$ denotes the number of completed future years lived by a person of particular age and gender at time T . In Australia we get a life expectancy of roughly 80.7 years for males and 84.9 for females born in 2013. Thus, comparing these numbers to a press release from October 2014 from the Australian Bureau of Statistics saying that ‘Aussie men now expected to live past 80’, we get similar results whereas our forecasts are slightly higher due to the consideration of mortality trends. It has to be noted that our results just show a modest increase in life expectancy (compared to the ABS press release) which is mainly due to a clear trend reduction, i.e., high values η , in mortality throughout the past ten years. If we do not believe in such a strong trend reduction, we may use a longer data history or a priori fix trend reduction parameters to a lower level. For example if we set $\zeta_{a,g} = 1972$, i.e., we assume trend reduction from the beginning of our data set, estimation gets more stable and parameter $\eta_{a,g}$ takes levels around 0.005 for males and 0.01 for females. In that case, due to negligible trend reduction, our model states that Australian men, born in 2013, are expected to live roughly 91 years and, thus, almost closing the gap to women who are also expected to live 91 years.

As a second example, we consider Austrian data from 1965 to 2014 with $t_0 = 2014$ and apply MCMC, again. In contrast to the Australian example, we assume $\kappa = 0$, i.e. no cohort effects, and set $\varepsilon_\alpha = \varepsilon_\beta = \varepsilon_\zeta = \varepsilon_\eta = 0$. Setting $\kappa = 0$ drastically reduces execution times, gives more stable estimates across time and, thus, this assumption is particularly useful for deriving long-term forecasts and life expectancy. Assuming $\varepsilon_\alpha = \varepsilon_\beta = \varepsilon_\zeta = \varepsilon_\eta = 0$ shows better fits of death probabilities across different ages. Regularisation parameters obtained by cross-validation are given by $c_\alpha = 150$, $c_\beta = 2 \cdot 10^6$, $c_\zeta = 10$ and $c_\eta = 2 \cdot 10^6$. The results, see fig:PDforecast4, show an overall improvement in mortality over all ages where the trend is particularly strong for young ages whereas the trend vanishes towards the age of 100, maybe implying a natural barrier for life expectancy. Due to sparse data the latter conclusion should be treated with the utmost caution. Furthermore, we see the classical hump of increased mortality around the age of 20. Estimates for $\zeta_{a,g}$ suggest that trend

acceleration switched to trend reduction throughout the past 20 to 30 years. Estimates for $\eta_{a,g}$ show that the speed of trend reduction is high, even stronger for males over most ages.




FIGURE 4.6. Parameter estimates for α, β, ζ and η .

Henceforth, based on forecasts for death probabilities, expected future life time can be estimated. To be consistent concerning longevity risk, mortality trends have to be included as a 60-year-old today will probably not have as good medication as a 60-year-old in several decades. In Austria we get a life expectancy of roughly 84.9 years for males and 86.9 for females born in 2014, i.e. significantly above official figures. For 60-year-old males and females we get expected future life times of 23.8 years and 26.1 years, respectively.

Using estimated values of death probabilities, weights and risk factor variances, as well as population forecasts (which are usually freely available at official statistical bureaus), our annuity model can be used to derive confidence bands for absolute numbers of deaths due to certain causes setting $Y_i = 1$. Using different MCMC parameter samples and using the parameter δ as described above, parameter risk and uncertainty in forecasts can be incorporated, respectively.

Mortality risk, longevity risk and Solvency II application. In light of the previous section, life tables can be projected into the future and, thus, it is straightforward to derive best estimate liabilities (BEL) of annuities and life insurance contracts. The possibility that death probabilities differ from an expected curve, i.e. estimated parameters do no longer reflect the best estimate and have to be changed, contributes to mortality or longevity risk, when risk is measured over a one year time horizon as in Solvency II and the duration of in-force insurance contracts exceeds this time horizon. In our model, this risk can be captured by considering various MCMC samples $(\hat{\theta}^h)_{h=1,\dots,m}$, yielding distributions of BELs. For example, taking $D(T, T+t)$ as the discount curve from time $T+t$ back to T and choosing an

MCMC sample $\hat{\theta}^h$ of parameters to calculate death probabilities $q_{a+t,g}^h(T+t)$ and survival probabilities $p_{a,g}^h(T)$ at age a with gender g , the BEL for a term life insurance contract which pays 1 unit at the end of the year of death within the contract term of d years is given by

$$A_{a,g}^T(\hat{\theta}^h) = D(T, T+1)q_{a,g}^h(T) + \sum_{t=1}^d D(T, T+t+1) \cdot {}_t p_{a,g}^h(T) q_{a+t,g}^h(T+t). \quad (4.3)$$

In a next step, this approach can be used as a building block for (partial) internal models to calculate basic solvency capital requirements (BSCR) for biometric underwriting risk under Solvency II, as illustrated in the following example.

Consider an insurance portfolio at time 0 with $n \in \mathbb{N}$ whole life insurance policies with lump sum payments $C_i > 0$, for $i = 1, \dots, n$, upon death at the end of the year. Assume that all assets are invested in an EU government bond (risk free under the standard model of the Solvency II directive) with maturity 1, nominal A_0 and coupon rate $c > -1$. Furthermore, assume that we are only considering mortality risk and ignore profit sharing, lapse, costs, reinsurance, deferred taxes, other assets and other liabilities, as well as the risk margin. Note that in this case, basic own funds, denoted by BOF_t , are given by market value of assets minus BEL at time t , respectively. Then, the BSCR at time 0 is given by the 99.5% quantile of the change in basic own funds over the period $[0, 1]$, denoted by ΔBOF_1 , which can be derived by, see (4.3),

$$\begin{aligned} \Delta\text{BOF}_1 &= \text{BOF}_0 - D(0, 1)\text{BOF}_1 = A_0(1 - D(0, 1)(1 + c)) - \sum_{i=1}^n C_i A_{a,g}^0(\hat{\theta}) \\ &\quad + \frac{D(0, 1)}{m} \sum_{h=1}^m \left(\sum_{i=1}^n C_i A_{a+1,g}^1(\hat{\theta}^h) + \sum_{i=1}^n \sum_{j=1}^{N_i^h} C_i (1 - A_{a+1,g}^1(\hat{\theta}^h)) \right). \end{aligned}$$

where $\hat{\theta} := \frac{1}{m} \sum_{h=1}^m \hat{\theta}^h$ and where N_1^h, \dots, N_n^h are independent and Poisson distributed with $\mathbb{E}[N_i^h] = q_{a_i, g_i}^h(0)$ with policyholder i belonging to age group a_i and of gender g_i . The distribution of the last sum above can be derived efficiently by Panjer recursion. This example does not require a consideration of market risk and it nicely illustrates how mortality risk splits into a part associated with statistical fluctuation (experience variance: Panjer recursion) and into a part with long-term impact (change in assumptions: MCMC). Note that by mixing N_i with common stochastic risk factors, we may include other biometric risks such as morbidity.

Consider a portfolio with 100 males and females at each age between 20 and 60 years, each having a 40-year term life insurance, issued in 2014, which provides a lump sum payment between 10 000 und 200 000 (randomly chosen for each policyholder) if death occurs within these 40 years. Using MCMC samples and estimates based on the Austrian data from 1965 to 2014 as given in the previous section, we may derive the change in basic own funds from 2014 to 2015 by (4.3). The distribution of change in BOFs is shown in Figure 4.7 where we observe a 99.5% quantile, i.e. the SCR, lying slightly above one million. If we did not consider parameter risk in the form of MCMC samples, the SCR would decrease by roughly 33%.

5. MODEL VALIDATION

Having estimated model parameters, it is straight-forward to derive validation techniques for our annuity model. For the first procedure, transform the data by (3.5) such that we may




FIGURE 4.7. Change in basic own funds, in thousand.

(very roughly) assume that sequence of number of deaths $(N'_{a,g,k}(t))_{t \in \{1, \dots, T\}}$ is i.i.d. over time. We then get explicit formulas for $\text{Var}(N'_{a,g,k})$, as well as $\text{Cov}(N'_{a,g,k}, N'_{a',g',k})$ and can then use samples from the Markov chains to derive quantiles. Then, these bounds can be compared to corresponding sample variances and sample covariances. In our Australian example 45.9 percent of all sample variances and covariances lie within five and 95 percent quantiles.

For the second procedure define

$$N^*_{a,g,k}(t) := \frac{N'_{a,g,k}(t) - \mathbb{E}[N'_{a,g,k}(t) | \Lambda_k(t)]}{\sqrt{\text{Var}(N'_{a,g,k}(t) | \Lambda_k(t))}} = \frac{N'_{a,g,k}(t) - m_{a,g} q_{a,g} w_{a,g,k} \Lambda_k(t)}{\sqrt{m_{a,g} q_{a,g} w_{a,g,k} \Lambda_k(t)}},$$

and note that the conditional central limit theorem implies $N^*_{a,g,k}(t) \rightarrow N(0, 1)$ in distribution as $m_{a,g}(t) \rightarrow \infty$ where $N(0, 1)$ denotes the standard normal distribution. Knowing that $\text{Cov}(N^*_{a,g,k}(t), N^*_{a',g',k'}(t)) = 0$ for all $k \neq k'$, we may as well test for correlation with a simple t -test. Applying this procedure to Australian data, we get that 88.9 percent of all tests are accepted at a five percent significance level.

A third possibility is to test for serial correlation in $(N^*_{a,g,k}(t))_{t \in \{1, \dots, T\}}$, e.g., via the Breusch–Godfrey test. Applying this validation procedure on Australian data gives that the null hypothesis, i.e., that there is no serial correlation of order $1, 2, \dots, 10$, is not rejected at a five percent level in 93.8 percent of all cases. Serial correlation is interesting insofar as there may be serial causalities between a reduction in deaths due to certain death causes and a possibly lagged increase in different ones, see Figure 1.1. Note that we already remove a lot of dependence via time-dependent weights and death probabilities.

Finally, we may use estimates for risk factor realisations to test whether risk factors are gamma distributed with mean one and variance σ_k^2 or not, e.g., via the Kolmogorov–Smirnov test. Note that estimates for risk factor realisations can either be obtained via MCMC based on the maximum a posteriori setting or by Equations (3.2) or (3.4). For Australia, this test gives acceptance of the null hypothesis for all risk factors on all suitable levels of significance.

For choosing a suitable family for mortality trends, information criteria such as AIC, BIC, or DIC can be applied straight away. The decision how many risk factors to use cannot be answered by traditional information criteria since a reduction in risk factors leads to a different data structure. It also depends on the ultimate goal. For example, if the development of all death causes is of interest, then a reduction of risk factors is not wanted.

On the contrary, in the context of annuity portfolios several risk factors may be merged to one risk factor as their contributions to the risk of the total portfolio are small.

6. CONCLUSION

Our approach provides a useful actuarial tool with numerous applications such as stochastic mortality modelling, P&L derivation in annuity and life insurance portfolios as well as (partial) internal model applications. Yet, there exists a fast and numerically stable algorithm to derive loss distributions exactly, even for large portfolios. We provide various estimation procedures based on publicly available data. The model allows for various other applications, including mortality forecasts. Compared to the Lee–Carter model, we have a more flexible framework, get tighter bounds and can directly extract several sources of uncertainty. Straightforward model validation techniques are available.

REFERENCES

- [1] A. J. G. Cairns, D. Blake, K. Dowd, G. D. Coughlan, D. Epstein, A. Ong, and I. Balevich, *A quantitative comparison of stochastic mortality models using data from England and Wales and the United States*, North American Actuarial Journal **13** (2009), no. 1, 1–35.
- [2] J. Hirz, U. Schmock, and Shevchenko P. V., *Modelling annuity portfolios and longevity risk with extended CreditRisk⁺*, preprint, arXiv: 1505.04757, 2016.
- [3] R. Kainhofer, M. Predota, and U. Schmock, *The new Austrian annuity valuation table AVÖ 2005R*, Mitteilungen der Aktuarvereinigung Österreichs **13** (2006), 55–135.
- [4] R. D. Lee and L. R. Carter, *Modeling and forecasting U.S. mortality*, Journal of the American Statistical Association **87** (1992), no. 419, 659–671.
- [5] U. Schmock, *Modelling Dependent Credit Risks with Extensions of Credit Risk⁺ and Application to Operational Risk*, <http://www.fam.tuwien.ac.at/~schmock/notes/ExtensionsCreditRiskPlus.pdf>, Lecture Notes, Version April 10, 2016.

(Jonas Hirz) DEPARTMENT OF FINANCIAL AND ACTUARIAL MATHEMATICS, TU WIEN, WIEDNER HAUPTSTR. 8/105-1, 1040 VIENNA, AUSTRIA

E-mail address: `hirz@fam.tuwien.ac.at`

(Uwe Schmock) DEPARTMENT OF FINANCIAL AND ACTUARIAL MATHEMATICS, TU WIEN, WIEDNER HAUPTSTRASSE 8-10, 1040 VIENNA, AUSTRIA

E-mail address: `schmock@fam.tuwien.ac.at`

URL: <http://www.fam.tuwien.ac.at/~schmock/>

(Pavel V. Shevchenko) APPLIED FINANCE AND ACTUARIAL STUDIES, MACQUARIE UNIVERSITY, NSW, AUSTRALIA

E-mail address: `pavel.shevchenko@mq.edu.au`